

Learning without Recall: A Case for Log-Linear Learning^{*}

Mohammad Amin Rahimian & Ali Jadbabaie^{*}

^{*} Department of Electrical and Systems Engineering, University of
Pennsylvania, Philadelphia, PA 19104-6228 USA
(e-mail: jadbabai@seas.upenn.edu)

Abstract: We analyze a model of learning and belief formation in networks in which agents follow Bayes rule yet they do not recall their history of past observations and cannot reason about how other agents' beliefs are formed. They do so by making rational inferences about their observations which include a sequence of independent and identically distributed private signals as well as the beliefs of their neighboring agents at each time. Fully rational agents would successively apply Bayes rule to the entire history of observations. This leads to forebodingly complex inferences due to lack of knowledge about the global network structure that causes those observations. To address these complexities, we consider a "*Learning without Recall*" model, which in addition to providing a tractable framework for analyzing the behavior of rational agents in social networks, can also provide a behavioral foundation for the variety of non-Bayesian update rules in the literature. We present the implications of various choices for time-varying priors of such agents and how this choice affects learning and its rate.

1. INTRODUCTION & BACKGROUND

Agents exchange beliefs in social networks to benefit from each other's opinions and private information in trying to learn an unknown state of the world. Rational agents in a social network would apply Bayes rule successively to their observations at each step, which include not only their private signals but also the beliefs communicated by their neighbors. However, such repeated applications of Bayes rule in networks become very complex, especially if the agents are unaware of the global network structure. This is due to the fact that the agents at each step should use their local data that is increasing with time, and make very complex inferences about possible signal structures leading to their observations. Indeed, tractable modeling and analysis of rational behavior in networks is an important problem in network economics and have attracted much attention, Acemoglu et al. (2011); Mueller-Frank (2013); Mossel et al. (2014).

To avoid the complexities of fully rational inference, a variety of non-Bayesian update rules have been proposed that rely on the seminal work of DeGroot (1974) in linear opinion pooling, where agents update their opinions to a convex combination of their neighbors' beliefs and the coefficients correspond to the level of confidence that each agent puts in each of her neighbors. More recently, Jadbabaie et al. (2012, 2013) consider a variation of this model for streaming observations, where in addition to the neighboring beliefs the agents also receive private signals. Other forms of non-Bayesian rules were studied by Bala and Goyal (1998) who consider a variation of observational learning in which agents observe the action and pay-offs of their neighbors and make rational inferences about these action/pay-off correspondence together with the

choices made by their neighbors, but ignore the fact that their neighbors are themselves learning from their own observations. More recently, Eyster and Rabin (2010) consider models of autarkic play where players at each generation observe their predecessor but naïvely think that any predecessor's action relies solely on that player's private information, thus ignoring the possibility that successive generations are learning from each other.

A chief contribution of this paper is establishing a behavioral foundation for the existing non-Bayesian updates in the literature where the belief update rule has a log-linear structure. This paper addresses the question of how one can limit the information requirement of a Bayesian update and still ensure consensus or learning for the agents. Some of the non-Bayesian update rules have the property that they resemble the replication of a first step of a Bayesian update from a common prior, and the aim here is to formalize such a setup. For instance DeMarzo et al. (2003) interpret the weights in the DeGroot model as those assigned initially by rational agents to the noisy opinions of their neighbors based on their perceived precision. However, by repeatedly applying the same weights over and over again, the agents ignore the need to update these weights with the increasing information.

To this end, we propose the so-called *Learning without Recall* model as a belief formation and update rule for *Rational but Memoryless* agents. We show how such a scheme can provide convergence and learning for all agents in a strongly connected social network so long as the truth is identifiable through the aggregate observations of the agents across entire network. This is of particular interest, when the agents cannot distinguish the truth based solely on their private observations, and yet together they learn.

^{*} This work was supported by ARO MURI W911NF-12-1-0509.

2. THE LEARNING WITHOUT RECALL MODEL

Notation. Throughout the paper, $\mathbb{N}_0 = \{0\} \cup \mathbb{N}$ is the set of positive integers and zero. Boldface letters denote random variables, the identity matrix is denoted by I , and $\|\cdot\|$ is a vector norm. We use Greek letters to denote certain variables of interest such as: agents beliefs (μ), log-ratio of beliefs (ϕ), log-likelihood ratio of signals (λ), and log-ratio of initial prior beliefs (ψ). We consider a network of n agents labeled by $[n] := \{1, 2, \dots, n\}$, that interact according to a directed graph $\mathcal{G} = ([n], \mathcal{E})$, where $\mathcal{E} \subset [n] \times [n]$ is the set of directed edges. Each agent is labeled by a unique element of the set $[n]$. $\mathcal{N}(i) = \{j \in [n]; (j, i) \in \mathcal{E}\}$ is the neighborhood of agent i , which is the set of all agents whose beliefs can be observed by agent i . We let $\deg(i) = |\mathcal{N}(i)|$ be the degree of node i corresponding to the number of agent i 's neighbors.

Signals and Environment. We denote by Θ the finite set of states of the world. Also, $\Delta\Theta$ represents the space of all probability measures on the set Θ . Each agent's goal is to decide amongst the finitely many possibilities in the state space Θ . A random variable θ is chosen randomly from Θ by nature and according to the probability measure $\nu(\cdot) \in \Delta\Theta$, which satisfies $\nu(\theta) > 0, \forall \theta \in \Theta$. For each agent i , there exists a finite signal space denoted by \mathcal{S}_i , and given θ , $\ell_i(\cdot | \theta)$ is a probability measure on \mathcal{S}_i , which is referred to as the *signal structure* or *likelihood function* of agent i . The private signals for agent i are generated according to $\ell_i(\cdot | \theta)$. Let Ω be an infinite product space encompassing the state θ and the private signals $\{\mathbf{s}_{i,t}, i \in [n], t \in \mathbb{N}_0\}$. Let $\mathbb{P}\{\cdot\}$ be the probability measure on Ω with the corresponding expectation operator $\mathbb{E}\{\cdot\}$, assigning probabilities consistently with the distribution $\nu(\cdot)$ and the likelihood functions $\ell_i(\cdot | \theta)$, $i \in [n]$. Conditioned on θ , the random variables $\{\mathbf{s}_{i,t}, i \in [n], t \in \mathbb{W}\}$ are independent. Consequently, the privately observed signals are independent and identically distributed over time and independent across the agents.

Beliefs. We let $\mu_{i,t}(\cdot)$ be a probability distribution on the set Θ representing the *opinion* or *belief* at time t of agent i about the realized value of θ . The goal is to study asymptotic learning, i.e. for each agent to learn the true realized value of θ asymptotically; we denote this truth by $\theta \in \Theta$. Hence, learning amounts to having $\mu_{i,t}(\cdot)$ converge to a point mass centered at θ , where the convergence could be in probability or in the stronger almost sure sense that we use in this work.

At $t = 0$ the value $\theta = \theta$ is selected by nature. Followed by that, $\mathbf{s}_{i,0}$ for each $i \in [n]$ is realized and observed by agent i . Then the agent forms an initial Bayesian opinion $\mu_{i,0}(\cdot)$ about the value of θ . Given $\mathbf{s}_{i,0}$, and using Bayes rule for each agent $i \in [n]$, the initial belief in terms of the observed signal $\mathbf{s}_{i,0}$ is given by:

$$\mu_{i,0}(\check{\theta}) = \frac{\nu_i(\check{\theta})\ell_i(\mathbf{s}_{i,0} | \check{\theta})}{\sum_{\hat{\theta} \in \Theta} \nu_i(\hat{\theta})\ell_i(\mathbf{s}_{i,0} | \hat{\theta})}, \quad (1)$$

where $\nu_i(\cdot) \in \Delta\Theta$ is an initial full-support ($\nu_i(\hat{\theta}) > 0, \forall \hat{\theta} \in \Theta$) prior belief associated with agent i ; it represents the subjective biases of agent i even before making any

observations. In particular, for an unbiased agent we have that $\nu(\hat{\theta}) = 1/|\Theta|, \forall \hat{\theta} \in \Theta$.

If we assume that each agent i knows the initial priors of her neighbors: $\nu_j(\cdot), j \in \mathcal{N}(i)$,¹ then Bayes rule can be exploited to derive the refined opinion $\mu_{i,1}(\cdot)$ after agent i observes her neighbors' initial beliefs $\{\mu_{j,0}(\cdot); j \in \mathcal{N}(i)\}$ given by (1). Following Rahimian et al. (2014) this leads to

$$\mu_{i,1}(\check{\theta}) = \frac{\nu_i(\check{\theta})\ell_i(\mathbf{s}_{i,0} | \check{\theta}) \left(\prod_{j \in \mathcal{N}(i)} \frac{\mu_{j,0}(\check{\theta})}{\nu_j(\check{\theta})} \right)}{\sum_{\hat{\theta} \in \Theta} \nu_i(\hat{\theta})\ell_i(\mathbf{s}_{i,0} | \hat{\theta}) \left(\prod_{j \in \mathcal{N}(i)} \frac{\mu_{j,0}(\hat{\theta})}{\nu_j(\hat{\theta})} \right)}. \quad (2)$$

How should the belief be updated in the next time steps? A Bayesian agent who cannot recall how the neighbors' beliefs are updated might instead imitate the preceding structure for all the following time steps $t > 1$. Indeed, in (2) we can substitute $\mathbf{s}_{i,t}$, $\mu_{i,t}(\cdot)$, and $\mu_{j,t-1}(\cdot)$ for $\mathbf{s}_{i,0}$, $\mu_{i,1}(\cdot)$, $\mu_{j,0}(\cdot)$. How should the agent update the priors? We study various choices for $\nu_j(\cdot), \forall j$.

To proceed, for each agent i we also replace $\nu_j(\cdot), \forall j$ with a time-varying distribution $\xi_{i,j}(\cdot, t)$, and argue that a *Rational but Memoryless* agent i would make her rational inference about the opinion $\mu_{j,t-1}(\cdot)$ that agent j reports to her at time t according to some time-varying prior $\xi_{i,j}(\cdot, t), j \in [n]$. Here, any choice of distributions $\xi_{i,j}(\cdot, t), j \in [n]$ should satisfy the information constraints of a *Rational but Memoryless* agent, so long as such a choice does not require an agent to recall any information other than what she has just observed $\mathbf{s}_{i,t}$ and what her neighbors have just reported to her $\mu_{j,t-1}(\cdot), j \in \mathcal{N}(i)$. A memoryless yet rational agent of this type can process the beliefs of her neighbors, but cannot recall how these beliefs were formed. Accordingly, (2) becomes

$$\mu_{i,t}(\check{\theta}) = \frac{\xi_{i,i}(\check{\theta}, t)\ell_i(\mathbf{s}_{i,t} | \check{\theta}) \left(\prod_{j \in \mathcal{N}(i)} \frac{\mu_{j,t-1}(\check{\theta})}{\xi_{i,j}(\check{\theta}, t)} \right)}{\sum_{\hat{\theta} \in \Theta} \xi_{i,i}(\hat{\theta}, t)\ell_i(\mathbf{s}_{i,t} | \hat{\theta}) \left(\prod_{j \in \mathcal{N}(i)} \frac{\mu_{j,t-1}(\hat{\theta})}{\xi_{i,j}(\hat{\theta}, t)} \right)}, \quad (3)$$

for all $\check{\theta} \in \Theta$ and at any $t > 1$. In writing (3), the time one update is regarded as a function that maps the priors, the private signal, and the neighbors' beliefs to the agent's posterior belief; and in using the time one update in the subsequent steps as in (3), every time agent i regards each of her neighbors $j \in \mathcal{N}(i)$ as having started from some prior belief $\xi_{i,j}(\cdot, t)$ and arrived at their currently reported belief $\mu_{j,t-1}(\cdot)$ directly after observing a private signal, hence rejecting any possibility of a past history. Such a rule is of course not the optimum Bayesian update of agent i 's belief at any step $t > 1$, because the agent is not taking into account the complete observed history of her private signals and neighboring beliefs and is instead, basing her inference entirely on the immediately observed signal and

¹ The assumption of (common) knowledge of priors in the case of rational (Bayesian) agents can be justified as follows: given the same observations of an agent j and in the absence of any past observations or additional data, any agent i should make the same (rational) inference; in the sense that starting from the same belief about the unknown, their updated beliefs given the same observations would be the same, or in Aumann's words, *rational agents cannot agree to disagree*, Aumann (1976).

neighboring beliefs; hence, the name *memoryless*. In the next section, we address the choice of random and time-varying priors $\xi_{i,j}(\cdot, t), j \in \mathcal{N}(i)$ while examining the properties of convergence and learning under the update rules in (3).

3. ASYMPTOTIC ANALYSIS OF CONVERGENCE & LEARNING WITH LOG-LINEAR UPDATES

We begin by forming the log-ratios of beliefs, signals, and priors under the true and false states as $\phi_{i,t}(\check{\theta}) := \log(\mu_{i,t}(\check{\theta})/\mu_{i,t}(\theta))$, $\lambda_{i,t}(\check{\theta}) := \log(\ell_i(s_{i,t}|\check{\theta})/\ell_i(s_{i,t}|\theta))$, and $\gamma_{i,j}(\check{\theta}, t) := \log(\xi_{i,j}(\check{\theta}, t)/\xi_{i,j}(\theta, t))$ for all i, j and t . Consequently, (3) can be linearized as follows:

$$\begin{aligned} \phi_{i,t}(\check{\theta}) &= \gamma_{i,i}(\check{\theta}, t) + \lambda_{i,t}(\check{\theta}) \\ &+ \sum_{j \in \mathcal{N}(i)} (\phi_{j,t-1}(\check{\theta}) - \gamma_{i,j}(\check{\theta}, t)). \end{aligned} \quad (4)$$

The network graph structure is encoded by its adjacency matrix A defined as $[A]_{ij} = 1 \iff (j, i) \in \mathcal{E}$, and $[A]_{ij} = 0$ otherwise. For a strongly connected \mathcal{G} the Perron-Frobenius theory, cf. (Seneta, 2006, Theorem 1.5), implies that A has a simple positive real eigenvalue, denoted by $\rho > 0$, which is equal to its spectral radius. Moreover, the left eigenspace associated with ρ is one-dimensional with the corresponding eigenvector $\bar{\alpha} = (\alpha_1, \dots, \alpha_n)^T$, uniquely satisfying $\sum_{i=1}^n \alpha_i = 1$, $\alpha_i > 0, \forall i \in [n]$, and $\bar{\alpha}^T A = \rho \bar{\alpha}^T$. The quantity α_i is also known as the eigenvector centrality of vertex i in the network, cf. (Newman, 2010, Section 7.2). Multiplying both sides of (4) by α_i and summing over all i we obtain that

$$\begin{aligned} \Phi_t(\check{\theta}) &= \text{tr}\{\Xi_t(\check{\theta})\} + \Lambda_t(\check{\theta}) \\ &+ \rho \Phi_{t-1}(\check{\theta}) - \text{tr}\{\Xi_t(\check{\theta}) A^T\} \\ &= \sum_{\tau=0}^t \rho^\tau (\Lambda_{t-\tau} + \text{tr}\{(I - A^T)\Xi_{t-\tau}(\check{\theta})\}) \end{aligned} \quad (5)$$

where $\Phi_t(\check{\theta}) := \sum_{i=1}^n \alpha_i \phi_{i,t}(\check{\theta})$ and $\Lambda_t(\check{\theta}) := \sum_{i=1}^n \alpha_i \lambda_{i,t}(\check{\theta})$ are global (network-wide) random variables, and $\Xi_t(\check{\theta})$ is a random $n \times n$ matrix whose i, j -th entry is given by $[\Xi_t(\check{\theta})]_{i,j} = \alpha_i \gamma_{i,j}(\check{\theta}, t)$. At each epoch of time, $\Phi_t(\check{\theta})$ characterizes how biased (away from the truth and towards $\check{\theta}$) the network beliefs are, and $\Lambda_t(\check{\theta})$ measures the information content of the received signal across all the agents in the network. Note that in writing (5), we use the fact that

$$\begin{aligned} \sum_{i=1}^n \alpha_i \sum_{j \in \mathcal{N}(i)} \phi_{j,t-1}(\check{\theta}) &= \bar{\alpha}^T A \bar{\phi}_{t-1}(\check{\theta}) \\ &= \rho \bar{\alpha}^T \bar{\phi}_{t-1}(\check{\theta}) = \rho \Phi_{t-1}(\check{\theta}), \end{aligned}$$

where $\bar{\phi}_t(\check{\theta}) := (\phi_{1,t}(\check{\theta}), \dots, \phi_{n,t}(\check{\theta}))^T$. On the other hand, since the received signal vectors $\{s_{i,t}, i \in [n], t \in \mathbb{N}_0\}$ are i.i.d. over time, $\{\Lambda_t(\check{\theta}), t \in \mathbb{N}_0\}$ constitutes a sequence of i.i.d. random variables satisfying

$$\mathbb{E}\{\Lambda_t(\check{\theta})\} = \sum_{i=1}^n \alpha_i \mathbb{E}\{\lambda_{i,t}(\check{\theta})\} = \sum_{i=1}^n \alpha_i \lambda_i(\check{\theta}) \leq 0, \quad (6)$$

where $\bar{\lambda}(\check{\theta}) := (\lambda_1(\check{\theta}), \dots, \lambda_n(\check{\theta}))^T :=$

$$-(D_{KL}(\ell_1(\cdot|\theta)||\ell_1(\cdot|\check{\theta})), \dots, D_{KL}(\ell_n(\cdot|\theta)||\ell_n(\cdot|\check{\theta})))^T,$$

and the non-positivity of (6) follows from the information inequality for the Kullback-Leibler divergence: $D_{KL}(\cdot||\cdot) \geq 0$, and is strict whenever $\ell_i(\cdot|\check{\theta}) \neq \ell_i(\cdot|\theta)$ for some i , i.e. $\exists s \in \mathcal{S}_i, i \in [n]$ such that $\ell_i(s|\check{\theta}) \neq \ell_i(s|\theta)$, cf. (Cover and Thomas, 2006, Theorem 2.6.3). In particular, if for all $\check{\theta} \neq \theta$ there exists an agent i with $\lambda_i(\check{\theta}) < 0$, then $\mathbb{E}\{\Lambda_t(\check{\theta})\} < 0$ and we say that the truth θ is *globally identifiable*. Indeed, if any agent is to learn the truth, then we need that $\Phi_t(\check{\theta}) \rightarrow -\infty$ as $t \rightarrow \infty$ for all the false states $\check{\theta} \neq \theta$. However, in a strongly connected graph every node has a degree greater than or equal to one so that $\rho \geq 1$, (Brualdi, 2011, Chapter 2). If $\rho > 1$, then the term $\rho^t \Lambda_0(\check{\theta})$ increases in variance as $t \rightarrow \infty$, and unless $\Lambda_0(\check{\theta}) < \epsilon$ with \mathbb{P} -probability one for some $\epsilon < 0$, almost sure convergence to $-\infty$ for $\Phi_t(\check{\theta})$ in (5) cannot hold true. In a directed circle where $\rho = 1$, one may take $\xi_{i,j}(\cdot, t) \equiv \nu_j(\cdot) \equiv \nu(\cdot)$ for all i, j and t . Consequently, at each epoch of time each agent i assumes that all her neighbors have started from the initial common prior belief $\nu(\cdot)$ and have arrived at their current beliefs directly, thus forgetting any history of observed signals and exchanged opinions. The geometric progression in (5) then reduces to sum of i.i.d. variables in \mathcal{L}^1 ; and by the strong law of large numbers, cf. (Feller, 1968, Section X.1), it converges almost surely to the mean value; hence, $\Phi_t(\check{\theta}) = \beta(\check{\theta}) + \sum_{\tau=0}^t \Lambda_\tau(\check{\theta}) \rightarrow \beta(\check{\theta}) + (t+1)\mathbb{E}\{\Lambda_0(\check{\theta})\} \rightarrow -\infty$, as $t \rightarrow \infty$, provided that $\mathbb{E}\{\Lambda_0(\check{\theta})\} < 0$, i.e. the truth is globally identifiable. Here $\beta(\check{\theta}) := \bar{\alpha}^T \bar{\psi}(\check{\theta})$, where $\bar{\psi}(\check{\theta}) := (\psi_1(\check{\theta}), \dots, \psi_n(\check{\theta}))^T$ is the stacked vector of initial prior log-ratios with $\psi_i(\check{\theta}) := \log(\nu_i(\check{\theta})/\nu_i(\theta)), \forall i \in [n]$. Thus as defined, $\beta(\check{\theta})$ measures the network-wide bias in the agents' initial priors toward the state $\check{\theta}$. Accordingly, (3) for a circular network with common priors becomes

$$\mu_{i,t}(\hat{\theta}) = \frac{\mu_{j,t-1}(\hat{\theta}) \ell_i(s_{i,t} | \hat{\theta})}{\sum_{\hat{\theta} \in \Theta} \mu_{j,t-1}(\hat{\theta}) \ell_i(s_{i,t} | \hat{\theta})}, \forall \hat{\theta} \in \Theta, \quad (7)$$

where $j \in [n]$ is the unique vertex $j \in \mathcal{N}(i)$. The above is the same as Bayesian update of a single agent (with no neighbors to communicate with), except that the self belief $\mu_{i,t-1}(\cdot)$ in the right-hand side is replaced by the belief $\mu_{j,t-1}(\cdot)$ of the unique neighbor $\{j\} = \mathcal{N}(i)$ in the circle.

Rahimian and Jadbabaie (2015) show that under (7) the agents in a directed circle learn the truth asymptotically exponentially fast and at the rate $\min_{\check{\theta} \neq \theta} (-1/n) \sum_{i=1}^n \lambda_i(\check{\theta})$, so that $\lim_{t \rightarrow \infty} \frac{1}{t} \Phi_t(\check{\theta}) = (1/n) \sum_{i=1}^n \lambda_i(\check{\theta})$, almost surely, for all $\check{\theta} \neq \theta$. Rahimian et al. (2015) show the application of the update rule in (7) to general strongly connected topologies (no necessarily circular), where agents have more than just a single neighbor in their neighborhoods. Accordingly, at every step of time agent i chooses a neighbor $j \in \mathcal{N}(i)$ independently at random and applies (7) with the reported belief of the chosen neighbor. The probabilities for the choice of neighbors at every point in time are given by a row stochastic matrix P , with entries $[P]_{ij} > 0$ for every $j \in \mathcal{N}(i)$. Each entry $[P]_{ij}$ is the probability for neighbor j being chosen by agent i at any point in time. Subsequently, we can show that the belief ratio for agent i at every time t is given as the sum of log-likelihood ratios of private signals of

various agents across the network and at times 0 to t . The choice of agent at every time is given by a random walk \mathbf{i}_τ , $\tau \in [t]$ that starts from node $\mathbf{i}_0 = i$ at time t and proceeds in the reversed time direction, eventually terminating at some node \mathbf{i}_t . Accordingly, the log-belief ratio of agent i at any time t can be expressed as $\phi_{i,t}(\theta) = \psi_{\mathbf{i}_t}(\theta) + \sum_{\tau=0}^t \lambda_{\mathbf{i}_\tau, t-\tau}(\theta)$. Next note from the ergodic theorem, cf. (Norris, 1999, Theorem 1.10.2), that the average time spent in any state $m \in [n]$ converges almost surely to its stationary probability π_m associated with the probability transition matrix P , and this together with the strong law yields $\lim_{t \rightarrow \infty} \frac{1}{t} \phi_{i,t}(\theta) = \sum_{i=1}^n \pi_i \lambda_i(\theta)$, almost surely for all $\theta \neq \theta$. Hence if the truth is globally identifiable, then in a strongly connected network every agent learns the truth at an asymptotic rate that is exponentially fast and is expressed above, as the sum of the relative entropies between the signal structures of every agent weighted by the stationary distribution of the random walk, which recovers the same asymptotic rate for the update proposed by Jadbabaie et al. (2013).

Fixing the priors over time will not result in convergence of beliefs, except in very specific cases as discussed above. In the sequel, we investigate the properties of convergence and learning under the update rules in (3), where the parameterizing priors $\xi_{i,j}(\cdot, t)$ are chosen to be random and time-varying variables, leading to the log-linear updating of the agents' beliefs over time. We distinguish two cases depending on whether the agents do not recall their own self-beliefs or they do, leading respectively to time-invariant or time-varying log-linear update rules.

3.1 Priors Set to a Geometric Average

It is notable that the memoryless Bayesian update in (3) has a log-linear structure similar to the Non-Bayesian update rules studied by Rahnama Rad and Tahbaz-Salehi (2010); Shahrampour and Jadbabaie (2013); Nedić et al. (2014); Lalitha et al. (2014); Bandyopadhyay and Chung (2014); and the roots for such a geometric averaging of the neighboring beliefs can be traced to logarithmic opinion pools as studied by Gilardoni and Clayton (1993) and Rufo et al. (2012). Motivated by this analogy, we propose setting the time-varying priors $\xi_{i,j}(\cdot, t)$, $j \in \mathcal{N}(i)$ of each agent i and at every time t , proportionally to the geometric average of the beliefs reported to her by all her neighbors at every time t : $\prod_{j \in \mathcal{N}(i)} \mu_{j,t-1}(\cdot)^{1/d(i)}$. Therefore, (3) becomes

$$\mu_{i,t}(\theta) = \frac{l_i(\mathbf{s}_{i,t} | \theta) \left(\prod_{j \in \mathcal{N}(i)} \mu_{j,t-1}(\theta) \right)^{1/d(i)}}{\sum_{\theta \in \Theta} l_i(\mathbf{s}_{i,t} | \theta) \left(\prod_{j \in \mathcal{N}(i)} \mu_{j,t-1}(\theta) \right)^{1/d(i)}}.$$

To analyze the evolution of beliefs with this choice of priors, let $\bar{\lambda}_t(\theta) := (\lambda_{1,t}(\theta), \dots, \lambda_{n,t}(\theta))^T$ be the stacked vector of log-likelihood ratios of received signals for all agents at time t . Hence, we can write the vectorized update $\bar{\phi}_t(\theta) = T \bar{\phi}_{t-1}(\theta) + \bar{\lambda}_t(\theta)$, where T is the normalized adjacency of the graph defined by $[T]_{ij} = \frac{1}{d(i)} [A]_{ij}$ for all i and j . We can now iterate the vectorized update to get $\bar{\phi}_t(\theta) = \sum_{\tau=0}^t T^\tau \bar{\lambda}_{t-\tau}(\theta) + T^t \bar{\psi}(\theta)$. Next note from the analysis of convergence for DeGroot model, cf. (Golub and Jackson, 2010, Proposition 1), that for a strongly connected network \mathcal{G} if it is aperiodic (meaning that one

is the greatest common divisor of the lengths of all its circles), then $\lim_{\tau \rightarrow \infty} T^\tau = \mathbf{1} \bar{s}^T$, where $\bar{s} := (s_1, \dots, s_n)^T$ is the unique left eigenvector associated with the unit eigenvalue of T and satisfying $\sum_{i=1}^n s_i = 1$, $s_i > 0$, $\forall i$. Hence, the Cesàro mean together with the strong law implies that $\lim_{t \rightarrow \infty} \frac{1}{t} \phi_{i,t}(\theta) = -\sum_{i=1}^n s_i \lambda_i(\theta)$, almost surely for all $\theta \neq \theta$, and the agents learn the truth asymptotically exponentially fast, at the rate $\min_{\theta \neq \theta} \sum_{i=1}^n -s_i \lambda_i(\theta)$.

3.2 Agents who Recall Their Self Beliefs

If the agents $i \in [n]$ recall their self-beliefs $\mu_{i,t-1}(\cdot)$, $i \in [n]$ when making decisions or performing inferences at time t , then we set $\xi_{i,i}(\cdot, t) \equiv \mu_{i,t-1}(\cdot)$ for all i and t . Furthermore, we set $\xi_{i,j}(\cdot, t) \equiv \mu_{j,t-1}(\cdot)^{\eta_t} / \zeta_j(t)$ for all $i, j \in \mathcal{N}(i)$ and t , where $\zeta_j(t) := \sum_{\theta \in \Theta} \mu_{j,t-1}(\theta)^{\eta_t}$ is the normalization constant to make the exponentiated probabilities sum to one. The choice of $0 < \eta_t < 1$ as time-varying exponents to be determined shortly, is motivated by the requirements of convergence under (3). Subsequently, we investigate the following log-linear update rule with time-varying coefficients

$$\phi_{i,t}(\theta) = \phi_{i,t-1}(\theta) + \lambda_{i,t}(\theta) + (1 - \eta_t) \sum_{j \in \mathcal{N}(i)} \phi_{j,t-1}(\theta),$$

where $1 - \eta_t$ is the weight that the agent puts on her neighboring beliefs (relative to her own) at any time t . Using $x_t := \rho(1 - \eta_t)$ and $B := (1/\rho)A$, the previous equation can be written in vectorized format as follows

$$\begin{aligned} \bar{\phi}_t(\theta) &= (I + x_t B) \bar{\phi}_{t-1}(\theta) + \bar{\lambda}_t(\theta) \\ &= \sum_{\tau=0}^t P^{(t,\tau)} \bar{\lambda}_\tau(\theta) + P^{(t,0)} \bar{\psi}(\theta), \end{aligned} \quad (8)$$

where $P^{(t,t)} := I$, and $P^{(t,\tau)} := \prod_{u=\tau+1}^t (I + x_u B)$ for $\tau < t$. Next note that $P^{(t,\tau)}$ can be expanded as follows

$$P^{(t,\tau)} = \sum_{j=0}^{t-\tau} M_j^{(t,\tau)} B^j,$$

where $M_0^{(t,\tau)} = 1$, $\tau \leq t$ and

$$M_j^{(t,\tau)} = \sum_{u_1=\tau+1}^{t-j+1} \sum_{u_2=u_1+1}^{t-j+2} \dots \sum_{u_j=u_{j-1}+1}^t x_{u_1} x_{u_2} \dots x_{u_j}.$$

Consequently, (8) can be rewritten as

$$\begin{aligned} \bar{\phi}_t(\theta) &= \\ &= \sum_{\tau=0}^t \sum_{j=0}^{t-\tau} M_j^{(t,\tau)} B^j \bar{\lambda}_\tau(\theta) + \sum_{j=0}^t M_j^{(t,0)} B^j \bar{\psi}(\theta), \end{aligned} \quad (9)$$

To proceed, for fixed τ and j let $M_j^{(\tau)} := \lim_{t \rightarrow \infty} M_j^{(t,\tau)}$. Next consider the summands in (9) for each τ , $0 \leq \tau \leq t$. Note that for j fixed, $\{B^j \bar{\lambda}_\tau(\theta), \tau \in \mathbb{N}_0\}$ is a sequence of independent and identically distributed random vectors. On the other hand, since B has unit spectral radius and given that $x_u > 0$, having $M_1^{(0)} := \sum_{u=1}^\infty x_u < \infty$ is sufficient to ensure that the random vectors $M_j^{(t,\tau)} B^j \bar{\lambda}_\tau(\theta)$ are all in \mathcal{L}^2 and have variances that are bounded uniformly in the choice of t, τ . This is because for any t, τ , and j we have that

$$M_j^{(t,\tau)} \leq M_j^{(\tau)} \leq M_j^{(0)} \leq \frac{1}{j!} \left(M_1^{(0)} \right)^j,$$

all as a consequence of positivity, $x_u > 0$. In particular, with $M_1^{(0)} < \infty$ we can bound

$$\left\| \sum_{j=0}^{\infty} M_j^{(0)} B^j \bar{\psi}(\check{\theta}) \right\| \leq \sum_{j=0}^{\infty} M_j^{(0)} \|B^j \bar{\psi}(\check{\theta})\| \leq \|\bar{\psi}(\check{\theta})\| \exp(M_1^{(0)}), \quad (10)$$

so that the contribution made by the initial bias of the network is asymptotically bounded and therefore sub-dominant when

$$\lim_{t \rightarrow \infty} \sum_{\tau=0}^t \sum_{j=0}^{t-\tau} M_j^{(t,\tau)} B^j \bar{\lambda}_\tau(\check{\theta}) = (-\infty)_n,$$

almost surely; here, by $(-\infty)_n$ we mean the entry-wise convergence of the column vector to $-\infty$ for the each of the n entries, corresponding to the n agents. In the sequel we investigate conditions under which this almost sure convergence would hold true.

We begin by noting that the condition $M_1^{(0)} < \infty$ is indeed necessary for convergence, because if $M_1^{(0)} = \infty$, then the term $M_1^{(t,0)} B \bar{\lambda}_0(\check{\theta})$ appearing in (9) for $j = 1$ and $\tau = 0$ increases unbounded in its variance as $t \rightarrow \infty$, so that (9) cannot converge in an almost sure sense.

Next note that with $M_1^{(0)} < \infty$ we can invoke Kolmogorov's criterion, (Feller, 1968, Section X.7), to get that as $t \rightarrow \infty$ the summation in (9) converges almost surely to its expected value, i.e. the following almost sure limit holds true

$$\lim_{t \rightarrow \infty} \phi_{i,t}(\check{\theta}) = \sum_{\tau=0}^{\infty} \sum_{j=0}^{\infty} M_j^{(\tau)} [B^j \bar{\lambda}(\check{\theta})]_i + \sum_{j=0}^{\infty} M_j^{(0)} [B^j \bar{\psi}(\check{\theta})]_i, \quad (11)$$

the second term being bounded per (10). Moreover, for a strongly connected social network \mathcal{G} , if it is aperiodic, then the matrix B is a primitive matrix; and in particular for all $j \geq d := \text{diam}(\mathcal{G}) + 1$, every entry of B^j is strictly greater than zero, and if the truth is globally identifiable then one can take an absolute constant $\epsilon \geq 0$ such that $[B^j \bar{\lambda}(\check{\theta})]_i < -\epsilon$ whenever $j \geq d$, while $[B^j \bar{\lambda}(\check{\theta})]_i \leq 0$ for any j . Subsequently, we get that

$$\sum_{\tau=0}^{\infty} \sum_{j=0}^{\infty} M_j^{(\tau)} [B^j \bar{\lambda}(\check{\theta})]_i \leq -\epsilon \sum_{\tau=0}^{\infty} \sum_{j=d}^{\infty} M_j^{(\tau)}. \quad (12)$$

Combining the results of (10), (11) and (12) leads to the following characterization: all agents will learn the truth (that is $\lim_{t \rightarrow \infty} \phi_{i,t}(\check{\theta}) = -\infty$, almost surely for all i and any $\check{\theta} \neq \theta$), if $M_1^{(0)} < \infty$ and $\sum_{\tau=0}^{\infty} \sum_{j=d}^{\infty} M_j^{(\tau)} = \infty$. Notice the preceding conditions are indeed not far from necessity. Firstly, we need $M_1^{(0)} < \infty$ to bound the growth of variance for convergence, as noted above. Moreover, with B having a unit spectral radius we can bound

$$|[B^j \bar{\lambda}(\check{\theta})]_i| \leq \|B^j \bar{\lambda}(\check{\theta})\| \leq \|\bar{\lambda}(\check{\theta})\|,$$

so that we can lower-bound $\lim_{t \rightarrow \infty} \phi_{i,t}(\check{\theta})$ in (11) as follows

$$-\|\bar{\lambda}(\check{\theta})\| \sum_{\tau=0}^{\infty} \sum_{j=0}^{\infty} M_j^{(\tau)} - \|\bar{\psi}(\check{\theta})\| \exp(M_1^{(0)}) \leq \lim_{t \rightarrow \infty} \phi_{i,t}(\check{\theta}).$$

Consequently, if $\sum_{\tau=0}^{\infty} \sum_{j=0}^{\infty} M_j^{(\tau)} < \infty$, then $\lim_{t \rightarrow \infty} \phi_{i,t}(\check{\theta})$ is almost surely bounded away from $-\infty$ and agents do not learn the truth.

For a strongly connected and aperiodic social network \mathcal{G} , matrix B has a single eigenvalue at one (corresponding to the largest eigenvalue of the adjacency A) and all of the other eigenvalues of B have magnitudes strictly less than one. Therefore, by the iterations of the power method, (Newman, 2010, Section 11.1), we know that $[B^j \bar{\lambda}(\check{\theta})]_i$ converges to $\sum_{k=1}^n \alpha_k \lambda_k(\check{\theta})$ for every i , and the convergence is geometrically fast in the magnitude-ratio of the first and second largest eigenvalues of the adjacency matrix A . Subsequently, for a strongly connected and aperiodic social network \mathcal{G} , we can replace $-\epsilon$ and d in (12) by $\epsilon + \sum_{i=1}^n \alpha_i \lambda_i(\check{\theta}) < 0$ and some constant $D(\epsilon)$. Here, $\epsilon > 0$ is a small but arbitrary and $D(\epsilon)$ is chosen large enough in accordance with the geometric rate of $[B^j \bar{\lambda}(\check{\theta})]_i \rightarrow \sum_{k=1}^n \alpha_k \lambda_k(\check{\theta})$, such that $|[B^j \bar{\lambda}(\check{\theta})]_i - \sum_{k=1}^n \alpha_k \lambda_k(\check{\theta})| < \epsilon$ for all $j \geq D(\epsilon)$, and the analysis of convergence and its rate can thus be refined.

Furthermore, having

$$K_1 < \liminf_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^t \sum_{j=d}^{t-\tau} M_j^{(t,\tau)} \leq \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^t \sum_{j=d}^{t-\tau} M_j^{(t,\tau)} < K_2$$

for some positive constants $0 < K_1 < K_2$ implies that the learning rate is asymptotically exponentially fast, as was the case for all the other update rules that we discussed in this paper. However, depending on how slow or fast (compared to t) is the convergence $\sum_{\tau=0}^t \sum_{j=d}^{t-\tau} M_j^{(t,\tau)} \rightarrow \infty$ as $t \rightarrow \infty$, the almost sure asymptotic rate at which for some $\check{\theta} \neq \theta$ and $i \in [n]$, $\mu_{i,t}(\check{\theta}) \rightarrow 0$ as $t \rightarrow \infty$ could be slower or faster than an exponential.

Time-Invariant Log-Linear Updates with Weighted Self-Beliefs.

For such agents as in Subsection 3.2 who recall their immediate self-beliefs $\mu_{i,t}(\cdot)$, if we relax the requirement that $\xi_{i,i}(\cdot, t) \equiv \mu_{i,t-1}(\cdot)$ then it is possible to achieve asymptotic almost sure exponentially fast learning using time-invariant updates just as in Subsection 3.1. In particular, for any $0 < \eta < 1$ fixed, we can set $\xi_{i,i}(\cdot, t)$ proportional to $\mu_{i,t-1}(\cdot)^\eta$ for all i , and we can further set $\xi_{i,j}(\cdot, t)$ at every time t , for any i , and all $j \in \mathcal{N}(i)$ proportional to $\mu_{j,t-1}(\cdot)^{1-(1-\eta)/d(i)}$. Subsequently, (3) becomes

$$\mu_{i,t}(\check{\theta}) = \frac{l_i(\mathbf{s}_{i,t} | \check{\theta}) \mu_{i,t-1}(\check{\theta})^\eta \left(\prod_{j \in \mathcal{N}(i)} \mu_{j,t-1}(\check{\theta}) \right)^{\frac{1-\eta}{d(i)}}}{\sum_{\hat{\theta} \in \Theta} l_i(\mathbf{s}_{i,t} | \hat{\theta}) \mu_{i,t-1}(\hat{\theta})^\eta \left(\prod_{j \in \mathcal{N}(i)} \mu_{j,t-1}(\hat{\theta}) \right)^{\frac{1-\eta}{d(i)}}}. \quad (13)$$

To analyze (13), we form the log-belief and likelihood ratios and set $B = (\eta I + (1 - \eta)T)$, where T is the same

normalized adjacency as in Subsection 3.1. Hence, we recover the vectorized iterations $\bar{\phi}_t(\bar{\theta}) = B\bar{\phi}_{t-1}(\bar{\theta}) + \bar{\lambda}_t(\bar{\theta}) = \sum_{\tau=0}^t B^\tau \bar{\lambda}_{t-\tau}(\bar{\theta}) + B^t \bar{\psi}(\bar{\theta})$ and it follows from (Kemeny and Snell, 1960, Theorems 5.1.1 and 5.1.2) that for a strongly connected network \mathcal{G} , $\lim_{\tau \rightarrow \infty} B^\tau = 1\bar{s}^T$, where $\bar{s} := (s_1, \dots, s_n)^T$ is the unique stationary distribution associated with the Markov chain whose probability transition matrix is T (or equivalently B); whence it follows from the Cesàro mean and the strong law that the agents learn the truth asymptotically exponentially fast, at the rate $\min_{\bar{\theta} \neq \theta} \sum_{i=1}^n -s_i \lambda_i(\bar{\theta})$, similar to Subsection 3.1. Note that here, unlike Subsection 3.1 but similarly to Rahimian et al. (2015), we only use properties of ergodic chains and existence and uniqueness of their stationary distributions; hence, relaxing the requirement for the social network to be aperiodic.

Concluding Remarks. Learning without recall is a model of belief aggregation in networks by replicating the rule that maps the initial priors, neighboring beliefs, and private signal to Bayesian posterior at one time step for all future time steps; this way the complexities of a fully rational inference at the forthcoming epochs are avoided, while some essential features of Bayesian inference are preserved. We showed how appropriate choice of priors for the so-called rational but memoryless agents in this model can provide a behavioral foundation for the log-linear structure of some of the non-Bayesian update rules that are studied in the literature. We further investigated the rate and requirements of asymptotic and almost sure learning with these choices of priors. In particular, we identified $\sum_{u=1}^{\infty} x_u < \infty$ as a necessary condition for convergence of beliefs with log-linear time-varying updates, in the case of agents who have no recollection of the past, excepting their own immediate beliefs. With x_t representing the relative weight on the neighboring beliefs, such agents facing individual identification problems can still learn the truth in a strongly connected and aperiodic network by relying on each other's observations; provided that as time evolves they put less and less weight on the neighboring beliefs and rely more on their private observations: $x_t \rightarrow 0$ as $t \rightarrow \infty$.

REFERENCES

- Acemoglu, D., Dahleh, M.A., Lobel, I., and Ozdaglar, A. (2011). Bayesian learning in social networks. *The Review of Economic Studies*, 78(4), 1201–1236.
- Aumann, R.J. (1976). Agreeing to disagree. *The annals of statistics*, 1236–1239.
- Bala, V. and Goyal, S. (1998). Learning from neighbours. *The Review of Economic Studies*, 65(3), 595–621.
- Bandyopadhyay, S. and Chung, S.J. (2014). Distributed estimation using bayesian consensus filtering. In *American Control Conference (ACC), 2014*, 634–641.
- Brualdi, R.A. (2011). *The mutually beneficial relationship of graphs and matrices*. 115. American Mathematical Soc.
- Cover, T. and Thomas, J. (2006). *Elements of Information Theory*. A Wiley-Interscience publication. Wiley.
- DeGroot, M.H. (1974). Reaching a consensus. *Journal of American Statistical Association*, 69, 118 – 121.
- DeMarzo, P.M., Vayanos, D., and Zwiebel, J. (2003). Persuasion bias, social influence, and unidimensional opinions. *The Quarterly Journal of Economics*, 118, 909–968.
- Eyster, E. and Rabin, M. (2010). Naive herding in rich-information settings. *American economic journal: microeconomics*, 2(4), 221–243.
- Feller, W. (1968). *An Introduction to Probability Theory and Its Applications*, volume 1. Wiley, 3rd edition.
- Gilardoni, G.L. and Clayton, M.K. (1993). On reaching a consensus using DeGroot's iterative pooling. *The Annals of Statistics*, 391–401.
- Golub, B. and Jackson, M.O. (2010). Naïve Learning in Social Networks and the Wisdom of Crowds. *American Economic Journal: Microeconomics*, 2(1), 112–149.
- Jadbabaie, A., Molavi, P., Sandroni, A., and Tahbaz-Salehi, A. (2012). Non-bayesian social learning. *Games and Economic Behavior*, 76(1), 210 – 225.
- Jadbabaie, A., Molavi, P., and Tahbaz-Salehi, A. (2013). Information heterogeneity and the speed of learning in social networks. *Revise and Resubmit, Review of Economic Studies*.
- Kemeny, J. and Snell, J. (1960). *Finite Markov Chains*. Van Nostrand, New York.
- Lalitha, A., Sarwate, A., and Javidi, T. (2014). Social learning and distributed hypothesis testing. In *2014 IEEE International Symposium on Information Theory (ISIT)*, 551–555.
- Mossel, E., Sly, A., and Tamuz, O. (2014). Asymptotic learning on bayesian social networks. *Probability Theory and Related Fields*, 158(1-2), 127–157.
- Mueller-Frank, M. (2013). A general framework for rational learning in social networks. *Theoretical Economics*, 8(1), 1–40.
- Nedić, A., Olshevsky, A., and Uribe, C.A. (2014). Nonasymptotic convergence rates for cooperative learning over time-varying directed graphs. *arXiv preprint arXiv:1410.1977*.
- Newman, M. (2010). *Networks: An Introduction*. Oxford University Press.
- Norris, J.R. (1999). *Markov Chains*. Cambridge University Press.
- Rahimian, M.A. and Jadbabaie, A. (2015). Learning without recall in directed circles and rooted trees. In *American Control Conference*, 4222–4227.
- Rahimian, M.A., Shahrampour, S., and Jadbabaie, A. (2015). Learning without recall by random walks on directed graphs. In *IEEE Conference on Decision and Control (CDC)*.
- Rahimian, M.A., Molavi, P., and Jadbabaie, A. (2014). (Non-) bayesian learning without recall. In *IEEE Conference on Decision and Control (CDC)*, 5730–5735.
- Rahnama Rad, K. and Tahbaz-Salehi, A. (2010). Distributed parameter estimation in networks. In *49th IEEE Conference on Decision and Control (CDC)*, 5050–5055. IEEE.
- Rufo, M., Martin, J., Pérez, C., et al. (2012). Log-linear pool to combine prior distributions: A suggestion for a calibration-based approach. *Bayesian Analysis*, 7(2), 411–438.
- Seneta, E. (2006). *Non-negative matrices and Markov chains*. Springer.
- Shahrampour, S. and Jadbabaie, A. (2013). Exponentially fast parameter estimation in networks using distributed dual averaging. In *52nd IEEE Conference on Decision and Control (CDC)*, 6196–6201.